

IN THE SPECIFICATION

At the top of the first page, just under the title, please insert: "This application is a continuation application of U.S. patent application serial No. 09/187,760, filed November 6, 1998, now U.S. Patent No. _____."

Please replace the following paragraphs in their entirety:

Page 8, second paragraph (page 8, lines 4-5):

FIG. 40 12 is a memory switch system according to another embodiment of the present invention.

Pages 13-14, paragraph spanning pages (page 13, lines 15-29 and page 14 lines 1-4):

In operation, a 64-byte packet 250 received at one of ports 210 by Port ASIC 245 is distributed (or divided) into four 16-byte packet portions (or other size portion, depending on the size of the packet 250 and the number of Memory ASICS) 250a, 250b, 250c, and 250d and stored in memory pool 225 via memory switch 220, as will be explained now in more detail with reference to FIG. 6. Switch interface 30 detects the arrival of a given packet 250 in one of the RX queues 22-1 ... 22-n associated with ports 210-1...210-n (step S10). Switch interface 30 determines, from address table 20, the range of memory addresses within memory pool 225 for storing the given packet 250 based on from which of the ports 210-1...210-n the given packet 250 arrived (step S20). Switch interface 30 immediately forwards a copy of the packet 250 (assuming it is the first packet in the frame, as determined in step S30, for example by determining whether it is the first packet received after a start of frame signal) to switch engine 230 (step S40). Alternatively, the switch engine 230 can independently receive a copy of the first packet in the frame by other instrumentalities. The first 64 bytes of an Ethernet frame will

include the frame header information that the switch engine 230 will use to determine how to forward the frame. Switch interface 30 also forwards a message to switch engine 230 that includes the memory pool 225 addresses at which the given packet 250 will be stored (step S50). Switch engine 230 will store this address in Table RAM 240, along with an identifier for the frame and then begin its packet forwarding determination operations.

Page 15, third full paragraph (page 15, lines 21-26):

If the packet that was stored was not the last packet in the frame (as determined in step S90, e.g. no signal has been received signaling the end of the frame), control will return to step S10 and the Port ASIC will continue to receive and store packets for the frame in memory pool 225. When the last packet is received, switch interface 30 will notify switch engine 230 to that effect (S100) and the **packet Port ASIC** will await further frames.

Pages 20-21, paragraph spanning pages (page 20 lines 24-29 and page 21, lines 1-6):

Moreover, it is possible to utilize the memory banks of DRAM devices so that, for example, a packet chunk is being received or stored in one memory bank of a given DRAM device, while another packet chunk is being retrieved from the another memory bank of the same given DRAM device or from another DRAM device in the memory pool 225. This full duplex operation capability thereby permits the system to achieve a higher bandwidth capability. In other words, the operations on different banks can overlap. While a given packet chunk is being read from one memory bank, other memory banks can be setting up new pages for the next packet chunk transfer. In contrast, if packet traffic is concentrated in only one memory bank, the bandwidth of a particular system will be more limited.

Pages 21-22, paragraph spanning pages (page 21, lines 7-29 and page 22, lines 1-2):

Commercially available memory devices are oriented towards PC applications whereby sequential data is transferred for long burst, e.g. 64 bytes or 128 bytes per transfer. Therefore, the page mode of DRAMs can be used efficiently when implementing the commercially available memory devices for PC applications. In the present invention, however, in which each 16-byte chunk is transferred to or from a different page, page mode is not as advantageous and the following modifications are essential. First, for a 64-bit wide DRAM implementation, a burst-size of "2" is used (i.e., in two (2) cycles, 16-bytes will be transferred to the DRAM). Second, the use of memory banks is maximized. Thus, data is spread across all available banks in the high bandwidth memory pool 225. By spreading data across available banks, the overlapping of memory operations ~~are~~ is allowed, thereby permitting a greater bandwidth to be achieved. A commercially available memory typically has about two (2) memory banks, while Dual Data Rate DRAMs can have up to about eight (8) memory banks. Third, additional pins are added to each of the DRAMs for use in DRAM signaling functions (such as "COMMAND" signals to the DRAMs. This permits a greater overlap to occur between COMMAND signals and DATA signals. FIG. 10A illustrates an example of memory pool 225 wherein each Memory ASIC interfaces with a memory device 227 having two memory banks. FIG. 10B illustrates an example of memory pool 225' wherein each Memory ASIC interfaces with a memory device 227' having four memory banks, which yields further overlapping of memory functions, and thus, greater bandwidth over the implementation in FIG. 10A.

Page 22, second full paragraph (page 22, lines 10-26):

Further advantages are obtained by utilizing Rambus memories to implement the memory pool. Some of the advantages of using Rambus DRAMs to implement the DRAM devices in the memory pool 225 are as follows. A Rambus DRAM (concurrent type) typically operates at about 600 MHz with 8 bits for data and a 31-pin interface. The peak per-pin-bandwidth is therefore (600MHz*8bits/31 pins) or 154 Mbits/s. In contrast, the bandwidth of conventional SDRAMs is limited by the SDRAM pin interface. The fastest commercially available SDRAM operates at about 143 MHz. For a 64-bit implementation, about 80 pins are required. Thus, the peak per-pin bandwidth of an SDRAM is therefore (125MHz*64bits/80 pins) or 100 Mbits/s. It should be noted that SDRAMs provide parallel synchronous **busses** for data transfers. Due to pinout constraints and cost constraints, the largest bus that exists on a single SDRAM chip is 32-bits in size. Multiple buses can be cascaded to form wider interfaces. However, since the ASIC pinout costs **increases** significantly with wider interfaces, practical limits today are buses of 64-bits, as factored in the above calculation.

///

///

///

///

///

///

///

///

///

///

///